

09 865066



APPLICATION FOR LETTERS PATENT OF THE UNITED STATES

CERTIFICATE OF MAILING "EXPRESS MAIL"	
"Express Mail"	
Mailing Label Number	
Date of Deposit	
I hereby certify that this paper or fee is being deposited with the United States Postal Service "EXPRESS MAIL POST OFFICE TO ADDRESSEE" Service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.	
(type or print name of person certifying)	
Signature	

SPECIFICATION

To all whom it may concern:

Be It Known, That we, **Clive Hoggart** and **James Griffin**, of London, United Kingdom and London, United Kingdom, respectively, have invented certain new and useful improvements in **METHOD AND APPARATUS FOR PREDICTING WHETHER A SPECIFIED EVENT WILL OCCUR AFTER A SPECIFIED TRIGGER EVENT HAS OCCURRED**, of which we declare the following to be a full, clear and exact description:



**METHOD AND APPARATUS FOR PREDICTING WHETHER A SPECIFIED
EVENT WILL OCCUR AFTER A SPECIFIED TRIGGER EVENT HAS OCCURRED**

Background of the Invention

5 This invention relates to a method and apparatus for predicting whether a specified event will occur for an entity after a specified trigger event has occurred for that entity. The invention is particularly related to, but in no way limited to, predicting customer behavior using a Bayesian statistical technique.

10 In many situations it is required to predict if and/or when an event will occur after a trigger. For example, businesses would like to predict if and when their customers are likely to leave after a particular event. The business is then able to take action to prevent loss of customers. Another case involves predicting if and when a bank customer is likely to take out a mortgage after a trigger such as a salary increase or change in marital status. The bank
15 would then be able to actively market its mortgages to specifically targeted groups of customers who are likely to be considering many different mortgage providers. Many other examples exist outside the banking and business fields. For example, predicting the time to death of patients after the trigger of a particular disease, which is known as “survival analysis” in the field of statistics.

20 Bayesian statistical techniques have been used to “learn” or make predictions on the basis of a historical data set. Bayes’ theorem is a fundamental tool for a learning process that allows one to answer questions such as “How likely is my hypothesis in view of these data?” For example, such a question could be “How likely is a particular future event to occur in view of these data?”

Bayes theorem is written as :

$$P(H / data) = \frac{P(data / H)P(H)}{P(data)}$$

Which can also be written as:

$$P(H / data) \propto P(data / H) \bullet P(H)$$

Because $P(data)$ is unconditional and thus does not depend on H .

The probability of H given the data, $P(H/data)$ is called the posterior probability of H .

- 5 The unconditional probability of H , $P(H)$ is called the prior probability of H and the probability of the data given H , $P(data/H)$ is called the likelihood of H . By using knowledge and experience about past data an assessment of the prior probability can be made. New data is then collected and used to update the prior probability following Bayes theorem to produce a posterior probability. This posterior probability is then a prediction in the sense that it is a
- 10 statement about the likelihood of a particular event occurring in the future. However, it is not simple to design and implement such Bayesian statistical methods in ways that are suited to particular practical applications.

Summary of the Invention

- 15 It is accordingly an object of the present invention to provide a method and apparatus for predicting whether a specified event will occur for an entity after a specified trigger event has occurred for that entity, which overcomes or at least mitigates one or more of the problems noted above.

- 20 According to an aspect of the present invention there is provided a method of predicting whether a specified event will occur for an entity after a specified trigger event has occurred for that entity, comprising the steps of:-

- accessing data about other entities for which the specified event has occurred in the past after the specified trigger event;
- accessing data about the entity for which the prediction is required;
- 25 • creating a Bayesian statistical model on the basis of at least the accessed data; and
- using the model to generate the prediction; wherein the data comprises a plurality of attributes associated with each entity and wherein creating the model comprises partitioning the attributes into a plurality of partitions.

A corresponding computer system is also provided for predicting whether a specified event will occur for an entity after a specified trigger event has occurred for that entity, comprising:-

- an input arranged to access data about other entities for which the specified event has occurred in the past after the specified trigger event; and wherein said input is further arranged to access data about the entity for which the prediction is required; wherein the data comprises a plurality of attributes associated with each entity;
- a processor arranged to create a Bayesian statistical model on the basis of at least the accessed data by partitioning the attributes into a plurality of partitions; and wherein the processor is further arranged to use the model to generate the prediction.

A corresponding computer program is provided, arranged to control a computer system in order to predict whether a specified event will occur for an entity after a specified trigger event has occurred for that entity, said computer program being arranged to control said computer system such that:-

- data is accessed about other entities for which the specified event has occurred in the past after the specified trigger event;
- data is accessed about the entity for which the prediction is required, wherein the data comprises a plurality of attributes associated with each entity;
- a Bayesian statistical model is created on the basis of at least the accessed data by partitioning the attributes into a plurality of partitions; and
- the model is used to generate the prediction.

This provides the advantage that it is possible to predict whether an event will occur after a trigger event. For example, the entities may be bank customers and using the method it is possible to predict whether a customer will leave a bank after having closed a loan with that bank. Data comprising customer attributes, such as the age, sex, salary, number of credit cards, number of loans, or current bank balance of the customers is used. A Bayesian statistical model is created and in doing this the attributes (which can be considered as existing in a space of attributes) are divided into a plurality of partitions. That is the space of

attributes is divided into partitions. By partitioning the attributes in this way the method is found to be particularly effective. Predictions are found to correspond well to empirical data in tests of the method as described further below and to give improved results as compared with prior art models which use global modeling techniques. By partitioning the attributes, the failings of global modeling techniques such as the method of Chen, Ibrahim and Sinha (see the section headed "references" below for bibliographic details of this publication) are avoided.

Preferably the Bayesian statistical model comprises a survival analysis type model which is arranged to take into account the assumption that the specified event will not occur for some of the entities. For example, in the case that the time to death of patients with a particular disease is being investigated, it is assumed that a proportion of these patients will not die and will be cured. Survival analysis models have previously used generalized linear models to account for customer/patient attributes. These global models typically lack sufficient flexibility to account for the variation across customers attributes in survival times. The present invention provides the advantage that a survival analysis model is adapted to fit a local model for customer attributes. An embodiment of the present invention maintains the proportional hazards property which although restrictive can be advantageous. The proportional hazards property implies that the ratio of the hazards for two customers is constant over time provided that their attributes do not change.

In another preferred embodiment the step of creating the model comprises fitting a Weibull distribution to the data within each partition. This provides the advantage that by fitting the Weibull distribution locally (i.e. within each partition) considerable modeling flexibility is gained. At the same time, the drawbacks of previous global survival models are overcome by using local modeling. This embodiment moves away from the restriction of proportional hazards.

Brief Description of the Drawings

Further benefits and advantages of the invention will become apparent from a consideration of the following detailed description given with reference to the accompanying drawings, which specify and show preferred embodiments of the invention.

5 Figure 1 is a flow diagram of a method for predicting whether a specified event will occur for an entity after a specified trigger event has occurred for that entity.

Figure 2 is a schematic diagram of a computer system for predicting whether a specified event will occur for an entity after a specified trigger event has occurred for that entity.

10 Figure 3 is a flow diagram of a method for predicting whether a specified event will occur for an entity after a specified trigger event has occurred for that entity.

Figure 4 is a flow diagram of another embodiment of a method for predicting whether a specified event will occur for an entity after a specified trigger event has occurred for that entity.

15 Figure 5 is a flow diagram of a method of sampling for a tessellation structure.

Figure 6 is a table containing example input data for the computer system of Figure 2 and example output data obtained from that computer system as well as corresponding empirical data.

Figure 7 is graph of the output data of Figure 6.

20

Detailed Description

Embodiments of the present invention are described below by way of example only. These examples represent the best ways of putting the invention into practice that are currently known to the Applicant although they are not the only ways in which this could be achieved.

25

Consider a business such as a bank. This bank may have beliefs, experience and past data about customer transactions. Using this information the bank can form an assessment of the prior probability that a particular customer will exhibit a certain behavior, such as leave

the bank. The bank may then collect new data about that customer's behavior and using Bayes' theorem can update the prior probability using the new observed data to give a posterior probability that the customer will exhibit the particular behavior such as leaving the bank. This posterior probability is a prediction in the sense that it is a statement of the likelihood of an event occurring. In this way the present invention uses Bayesian statistical techniques to make predictions about customer behavior. However, as mentioned above, it is not simple to design and implement such methods in ways that are suited to particular applications. The present invention involves such a method and is described in more detail below.

Figure 1 is a flow diagram of a method for predicting whether a specified event will occur for an entity after a specified trigger event has occurred for that entity. Data is accessed about entities for which a specified event has occurred in the past after a specified trigger event (see box 10 of Figure 1). The entities may be customers, individuals, or any other suitable item such as a computer system. For example, the data comprises customer attributes such as age, sex and salary for customers who have closed a loan and then left the bank. More data is then accessed (see box 11 of Figure 1) about an entity for which it is required to make a prediction. For example, this data may comprise customer attributes associated with customers for whom it is required to predict whether they will leave a bank after closing a loan.

A Bayesian statistical model is then created (see box 12 of Figure 1) on the basis of at least the accessed data and this model is used to generate the predictions. The process of generating the model comprises partitioning the attributes in to a plurality of partitions.

Two embodiments of the method of Figure 1 are now described. The first embodiment takes a Bayesian survival model and adapts it such that attribute data are partitioned. The second embodiment involves fitting a Weibull distribution to the customer attribute data within each partition. Both embodiments are described below with respect to a particular application, that of predicting if and/or when a customer will leave a bank after having paid off a loan. However, this embodiment is also suitable for other applications in

which it is required to predict whether a specified event will occur for an entity after a specified trigger event has occurred for that entity.

The methods of both these embodiments may be implemented using any suitable programming language executed on any suitable computing platform. For example, Matlab (trade mark) may be used together with a personal computer. A user interface is provided such as a graphical user interface to allow an operator to control the computer program, for example, to adjust the model, to display the results and to manage input of customer data. Any suitable form of user interface may be used as is known in the art.

Figure 2 is a schematic diagram of a computer system for predicting whether a specified event will occur for an entity after a specified trigger event has occurred for that entity. The computer system comprises a processor 23 which may be any suitable type of computing platform such as a personal computer or a workstation. The computer system has an input 25 which is arranged to receive data 21 about entities for which a specified event has occurred in the past after a specified trigger event. This input 25 is also arranged to receive data about an entity (or entities) for which it is required to predict if a specified event will occur after a specified trigger event has occurred. Using this data, which comprises a plurality of attributes associated with each entity, the processor generates a Bayesian statistical model and partitions the attributes into a plurality of partitions. Once the model is formed it is used by the processor 23 to generate predictions 24 about if and/or when the specified event will occur after the specified trigger event for one or more entities.

The first embodiment is now described:

A common problem faced by banks is customer attrition. In order to deal with this problem banks required the answer to the question "will customer A leave the bank?" We are interested in the case where customer attrition occurs after a particular event. For example, customers may leave a bank after having paid off a loan. If we can predict who will leave and the time between closing the account and leaving the bank, then action can be taken to prevent the customer leaving.

This problem is similar to the statistical subject of survival analysis. In a typical medical survival analysis problem the time to death of a patient with a particular disease is investigated. Typical models assume that all patients will eventually die from the disease. However, in the present invention it is assumed that a proportion of the customers will not leave the bank due to the particular event. In medicine this is equivalent to a proportion of the patients being cured and models which have accounted for this allow for a so called “cure rate”.

A Bayesian survival model has been developed (Chen, Ibrahim and Sinha, Journal of the American Statistical Association, 1999) which allows for a cure rate. The model described in the paper allows the cure rate to vary for individuals with different attributes by using a generalized linear model. A generalized linear model is a global model. In a global model an assumption is made about how the data is distributed as a whole and so global modeling is a search for global trends. However, all customers may not follow a global trend; some subpopulations of customers may differ radically from others. The present invention extends the work of Chen, Ibrahim and Sinha (1999) to model the customer attributes locally avoiding the failings of the global generalized linear model.

The first embodiment is now described with reference to Figure 3.

In order to create the Bayesian statistical model, first prior distributions are chosen on the basis of beliefs, experience and past data about customer attributes and behavior (see box 31 of Figure 3). For example, the prior distributions may be specified as gamma distributions. A tessellation structure and parameters for the model are then initialized (see box 32 of Figure 3) for example, by assigning random values. The customer attributes are considered as being represented in a customer attribute space and the tessellation structure represents division of this space into partitions.

Any suitable sampling method such as a Gibbs sampling method is then used to form a posterior probability distribution from the prior distributions and customer data. This is represented by box 40 of Figure 3. This process comprises sampling for the tessellation structure (box 33 of Figure 3) and sampling for a cure rate within each partition (box 34) by

making a standard draw from a gamma distribution (in the case that the prior distributions are modeled as gamma distributions). As well as this, the method comprises, for each customer, sampling for N , which is the number of latent risks (box 35). The number of latent risks is an indication of how likely a customer is to leave the bank. The greater the number of latent

5 risks the more likely the customer is to leave. In one example, sampling for N is achieved by making a standard draw from a Poisson distribution. The next stage involves sampling for parameters of the distribution of the latent risks. In one example, this is achieved by making standard draws for the parameters of a Weibull distribution.

The sampling steps of box 40 of Figure 3 are repeated until sufficient samples are

10 obtained to enable the posterior probability distribution to be described and “reconstructed”. For example, this is done by repeating the sampling steps for a pre-specified large number of iterations and assuming that sufficient samples will have been drawn (for example several thousand iterations). The results may then be compared with empirical data and the effect of further iterations assessed. Once sufficient samples have been obtained the model is said to

15 have converged. Thus in Figure 3 a decision point 37 is shown with the test “Has Markov chain converged?”. If the answer to this question is “no” and insufficient samples have been drawn the sampling method is repeated starting from box 33. If the answer to this question is “yes” then the posterior probability distribution is assumed to have been adequately described. In that case, the sampling method is repeated in order to draw samples from the reconstructed

20 probability distribution (box 38) and these samples are used to generate probabilities as to if and when each customer will leave the bank (box 39).

The step of sampling for the tessellation structure (box 33 of Figure 3) is shown in more detail in Figure 5. This is an iterative process which involves adjusting the tessellation structure if a parameter u is greater than a calculated acceptance ratio where u is a uniform

25 random variable between 0 and 1. The first step involves either adding a new hyperplane, removing an existing hyperplane or moving an existing hyperplane. Once this has been done a representation of the tessellation structure is revised in order to take into account the change. For example, the tessellation structure may be represented using a temporary hash table which

is recalculated to take into account the change (box 52). A marginal likelihood is then calculated (this is described in more detail below) (box 53) and an acceptance ratio also calculated (box 54). The parameter u is then uniformly drawn (box 55) using a sampling method. If u is greater than the acceptance ratio then no changes are made to the tessellation structure (box 58). However, if u is less than the acceptance ratio then the process is repeated (box 57).

The first embodiment and the way in which this extends the work of Chen, Ibrahim and Sinha is now described in more detail:

The approach described by Chen, Ibrahim and Sinha models the unknown number of cancerous cells, or more generally “risks”, in a patient. If a patient has no cancerous cells the patient is said to be cured, otherwise the risk is assumed to increase with the number of cancerous cells. The number of risks, denoted by N , is modeled as a Poisson distribution. The time to death due to risk i is denoted by Z_i . The model assumes that the random variables Z_1, \dots, Z_n are independent and identically distributed (i.i.d.) with a common distribution function $F(t) = 1 - S(t)$, where $S(t)$ is known as the survival function and represents the probability of surviving to time t . The overall survival function is given by the probability of surviving N risks until time t . This is written as

$$\begin{aligned} S_p(t) &= P(\text{alive at time } t) \\ &= P(N = 0) + P(Z_1 > t, \dots, Z_N > t, N \geq 1) \\ &= \exp(-\theta) + \sum_{k=1}^{\infty} S(t)^k \frac{\theta^k}{k!} \exp(-\theta) \\ &= \exp(-\theta + \theta S(t)) = \exp(-\theta F(t)) \end{aligned}$$

t is the response of interest, for example the time between a customer closing a loan and leaving the bank. The distribution function $F(t)$ of the risks Z can take any form, for example the Weibull distribution is used. However, it is not essential to use the Weibull distribution; any other suitable distribution can be used. The Weibull distribution has the following density function

$$p(t|\alpha, \lambda) = \lambda \alpha t^{\alpha-1} \exp(-\lambda t^\alpha)$$

Chen, Ibrahim and Sinha model the parameter of the Poisson distribution with a generalized linear model, thus

$$\theta = \exp(X' \beta),$$

a generalized linear model. A customer's attributes are denoted by X and β denotes the

5 parameters. Thus if we have p customer attributes X_1, \dots, X_p we will have parameters

β_1, \dots, β_p . This is a global model because the parameters, β , take the same value for each

customer. The unknown parameters of the model are $N_1, \dots, N_n, \lambda, \gamma$ and β where λ and γ

are the parameters of the Weibull distribution. As with most Bayesian models, the posterior distribution of the unknown parameters cannot be expressed analytically. The Gibbs sampler

10 is a widely used method for drawing random values from posterior distributions. The posterior distribution is reconstructed from the samples generated by the Gibbs sampler. To implement

a Gibbs sampler the full conditional distributions of the parameters are required. Sampling for

β is not standard. An algorithm exists to draw from the full conditional distribution of each

component of β . However the algorithm is relatively computationally expensive and p

15 draws will be required from it for each sweep of the Gibbs sampler.

Global models, such as that described by Chen, Ibrahim and Sinha are not always appropriate, particularly for a large set of customers. In that case a local model as described

in the present invention has been found to be more effective. The local model of the present invention is simple and more flexible than the generalized linear model used previously. The

20 space of customer attributes is split into disjoint sub-populations or partitions. The partitions are defined geometrically. For example, hyperplanes are used to divide the space of customer attributes. Within each sub-population a constant response θ is fit, the most simple of local models.

The unknown parameters of the model are $N_1, \dots, N_n, \alpha, \lambda, T$ and $\theta_1, \dots, \theta_m$ where T

25 denotes the tessellation structure with m sub-populations or partitions. We denote the

response in the partition j by θ_j , the number of observations in partition j by n_j , the latent

variables in partition j by N_{1j}, \dots, N_{n_jj} and the observations in partition j by t_{1j}, \dots, t_{n_jj} . A

Gibbs sampler (or any other suitable type of sampling method) is used to draw from the posterior distribution of the unknown parameters which is given by

$$p(\alpha, \lambda, N_1, \dots, N_n, \theta_1, \dots, \theta_m, T | t_1, \dots, t_n) \propto p(\alpha) p(\lambda) \prod_{j=1}^m p(\theta_j) \prod_{i=1}^{n_j} p(t_{ij} | N_{ij}, \alpha, \lambda) p(N_{ij} | \theta_j)$$

$$= p(\alpha) p(\lambda) \prod_{j=1}^m p(\theta_j) \exp \left\{ -\lambda \sum_{i=1}^{n_j} N_i t_{ij}^{\alpha_j - 1} \right\} \prod_{i=1}^{n_j} (N_i \lambda \alpha_i^{\alpha_i - 1})^{\theta_j} \frac{\theta_j^{N_{ij}} \exp(-\theta_j)}{N_{ij}!}$$

The following prior distributions are assigned

$$p(\theta_j) = Ga(\varphi_0, \varphi_1)$$

$$5 \quad p(\lambda) = Ga(\lambda_0, \lambda_1)$$

$$p(\alpha) = Ga(\alpha_0, \alpha_1)$$

which are all gamma distributions. However, it is not essential to use Gamma distributions to model the prior distributions. Any other suitable type of distribution can be used.

The Gibbs sampler (or other sampling method) draws from the following full conditional distributions

$$p(\alpha | \dots) \propto \alpha^{n + \alpha_0 - 1} \left(\prod_{i=1}^n t_i \right)^{\alpha} \exp \left\{ -\alpha \alpha_0 - \lambda \sum_{i=1}^n N_i t_i^{\alpha} \right\}$$

$$10 \quad p(\lambda | \dots) = Ga(\lambda | n + \lambda_0, \lambda_1 + \sum N_i t_i^{\alpha})$$

$$p(N_{ij} | \dots) = Pn(\theta_j \exp(-\lambda t_i^{\alpha})), \quad i = 1, \dots, n_j, j = 1, \dots, m$$

$$p(\theta_j, T | \dots) = p(T | \dots) p(\theta_j | T, \dots), \quad j = 1, \dots, m$$

where

$$p(\theta_j | T, \dots) = Ga(\varphi_0 + n_j, \varphi_1 + \sum_{i=1}^{n_j} N_{ij})$$

$$p(T | \dots) \propto p(N_1, \dots, N_n | T) p(T)$$

$$= p(T) \prod_{j=1}^m p(N_{1j}, \dots, N_{n_jj} | T)$$

Ga denotes the gamma distribution and Pn denotes the Poisson distribution. The example discussed here uses Poisson distributions to model the full conditional distributions, however,

any other suitable type of distribution can be used. An advantage of choosing the Poisson distribution is that marginal likelihoods are straightforward to calculate as described below.

To fit a local model the marginal likelihood $p(N_1, \dots, N_n)$ is required. The marginal likelihood is the likelihood of the data with the parameters θ integrated out.

- 5 The marginal likelihood is straightforward to evaluate in this model due to the nature of the Poisson distribution. If we assign θ a $\text{Gamma}(\theta_0, \theta_1)$ prior the marginal likelihood of the number of risks of each customer N_1, \dots, N_n is given by

$$\begin{aligned} p(N_1, \dots, N_n | \varphi_0, \varphi_1) &= \int \prod_{i=1}^n p(N_i | \theta) p(\theta | \varphi_0, \varphi_1) d\theta \\ &= \int \prod_{i=1}^n \frac{\theta^{N_i} \exp(-\theta)}{N_i!} \frac{\Gamma(\varphi_0)}{\varphi_1^{\varphi_0}} \theta^{\varphi_0-1} \exp(-\varphi_1 \theta) d\theta \\ &= \frac{\varphi_1^{\varphi_0}}{\Gamma(\varphi_0) \Pi(N_i!)} \int \theta^{\sum N_i + \varphi_0 - 1} \exp(-\theta(n + \varphi_1)) d\theta \\ &= \frac{\Gamma(\theta_0 + \sum N_i)}{(\varphi_1 + n)^{\varphi_0 + \sum N_i} \Pi(N_i!)} \frac{\varphi_1^{\varphi_0}}{\Gamma(\varphi_0)} \end{aligned}$$

- 10 Given the marginal distribution, the tessellation structure is sampled for using a Metropolis random walk, within the Gibbs sampler (or other sampling method).

The resulting sampler is computationally more efficient than the equivalent sampler for the generalized linear model described above. Sampling for β has been replaced by sampling for the tessellation structure and the responses within each partition, both of which are straightforward.

- 15 The method described above has been implemented using a computer system such as that illustrated in Figure 2. Figure 6 is a table containing example input data for the computer system of Figure 2 and example output data obtained from that computer system (using the method described immediately above) as well as corresponding empirical data. The first four columns 60 of the table in Figure 6 are headed “co-variates” and contain attribute values.
- 20 Each row of the table represents data for an individual bank customer. Columns 61 to 63 contain probability values which have either been obtained from empirical data (column 63), or which have been obtained from the method of the present invention (column 62), or from

the prior art method of Chen, Ibrahim and Sinha (column 61). The final column 64 of table 6 shows the number of observations that were available for each customer.

The probability values produced by the method of the present invention are closer to the empirical values than those produced by the prior art method of Chen, Ibrahim and Sinha.

- 5 For example, for the first customer whose data is contained in the first row of the table, the empirical probability value is 0.2795 and the probability value predicted using the method of the present invention is 0.2047 whereas the prior art method gave 0.4213.

Figure 7 shows a graph formed using the data of Figure 6 together with further data for other customers. The graph is a plot of the proportion of customers who are still with the
10 bank (or predicted to be still with the bank) against time in days. The results of the prior art Chen, Ibrahim and Sinha model are represented by the upper curve 71 and the results of the method of the present invention by the lower curve 72. A single point 73 is shown which indicates the proportion of customers still with the bank after 1 year. This data point is obtained from empirical data.

- 15 The data shown in Figures 6 and 7 which are produced from the method of the present invention are slight underestimates of the empirical data. This is because not all people who will leave the bank have actually left by the end of the experiment. This means that the actual proportion (from empirical data) of people who are still with the bank will be lower than predicted using the method of the present invention. Taking this into account, the predictions
20 of the present invention are actually even closer to the empirical data in Figure 7.

The second embodiment is now described with reference to Figure 4. As for the first embodiment, prior distributions are chosen (box 41) and the tessellation structure and parameters are initialized (box 42). Using the prior distributions and input customer data a Gibbs sampling method (or any other suitable sampling method) is then used to draw samples
25 in order to “reconstruct” the posterior probability distribution. This involves sampling for the tessellation structure (box 43) and then sampling for the parameters of the distribution of latent risks (box 44). This comprises taking standard draws for the parameters of the Weibull distribution (box 44). The next stage (box 45) comprises for each customer, sampling for N,

the number of latent risks. This is achieved by taking a standard draw from a Poisson distribution (or any other suitable distribution).

As in the first embodiment the sampling process is iterated until the posterior probability distribution has been adequately “reconstructed” (see box 46). This is achieved in
5 any of the ways described above for the first embodiment.

Once convergence has been achieved, the posterior probability distribution is assumed to be adequately “reconstructed” and samples are then drawn from it (box 47) using the sampling method of box 49. The samples drawn from the posterior probability distribution are then used to generate probabilities as to if and when each customer will leave the bank
10 (box 48).

The second embodiment is now described in more detail:

The second embodiment uses a local model and splits the space of customer attributes into disjoint sub-populations or partitions. The partitions are defined geometrically. For example, hyperplanes can be used to divide the space of customer attributes. Within each partition a
15 Weibull distribution is fitted which has the following density function:

$$p(t|\alpha, \lambda) = \lambda \alpha t^{\alpha-1} \exp(-\lambda t^\alpha)$$

In survival analysis t refers to the time of death of a patient. In a banking context t represents for example, the time between a customer closing a loan and leaving the bank.

The local Weibull distribution makes use of the following mixture representation of
20 the Weibull distribution:

$$p(t|u, \alpha) = \alpha u^{-1} t^{\alpha-1} I(t^\alpha < u)$$

$$p(u|\lambda) = \lambda^2 u \exp(-u\lambda)$$

as described by Walker and Gutierrez-Pera (see the section headed “references” below for bibliographic details).

It is straightforward to show that this mixture yields the marginal distribution

25 $p(t|\alpha, \lambda) = \lambda \alpha t^{\alpha-1} \exp(-\lambda t^\alpha)$

which is Weibull(α, λ).

The unknown parameters of the model are $u_1, \dots, u_n, \alpha_1, \dots, \alpha_m, \lambda_1, \dots, \lambda_m$ and the tessellation structure T with m sub-populations or partitions. The parameters of the Weibull distribution in partition j are denoted by α_j, λ_j , the number of observations in partition j is denoted by n_j and the latent variables in partition j are denoted by u_{1j}, \dots, u_{n_jj} , similarly we
 5 denote the observations in partition j by t_{1j}, \dots, t_{n_jj} . The posterior distribution of the unknown parameters is

$$p(\alpha_1, \dots, \alpha_m, \lambda_1, \dots, \lambda_m, u_1, \dots, u_n, T | t_1, \dots, t_n) = \prod_{j=1}^m p(\alpha_j) p(\lambda_j) \prod_{i=1}^{n_j} p(t_{ij} | u_{ij}, \alpha_j) p(u_{ij} | \lambda_j)$$

$$= \prod_{j=1}^m p(\alpha_j) p(\lambda_j) \prod_{i=1}^{n_j} \alpha \lambda^2 t_{ij}^{\alpha-1} \exp(-u_{ij} \lambda_j) I(t_{ij}^{\alpha_j} < u_{ij})$$

We take the following prior distributions for α and λ

$$p(\lambda_j) = Ga(\lambda_0, \lambda_1)$$

$$p(\alpha_j) = Ga(\alpha_0, \alpha_1)$$

10 However, it is not essential to represent the prior distributions using Gamma distributions. Any other suitable distributions can be used.

As with most Bayesian models, the posterior distribution of the unknown parameters cannot be expressed analytically. The Gibbs sampler (or any other suitable sampling method) is therefore used to draw random values from the posterior distribution. The posterior
 15 distribution is then reconstructed from the samples generated by the Gibbs (or other) sampler. To implement the Gibbs (or other) sampler the full conditional distributions of the parameters are required. In the present embodiment we draw from the following full conditional distribution

$$p(\alpha_1, \dots, \alpha_m, \lambda_1, \dots, \lambda_m, T | t_1, \dots, t_n, u_1, \dots, u_n)$$

$$= p(\alpha_1, \dots, \alpha_m, \lambda_1, \dots, \lambda_m | T, t_1, \dots, t_n, u_1, \dots, u_n) p(T | t_1, \dots, t_n, u_1, \dots, u_n)$$

$$p(u_1, \dots, u_n | \alpha_1, \dots, \alpha_m, \lambda_1, \dots, \lambda_m, T, t_1, \dots, t_n)$$

Given a tessellation structure $\alpha_1, \dots, \alpha_m, \lambda_1, \dots, \lambda_m$ and u_1, \dots, u_n are independent and their full conditional distributions are as follows:

$$p(\alpha_j | \dots) \propto \alpha_j^{n_j + \alpha_0 - 1} \exp \left\{ -\alpha_j \left(\alpha_1 - \sum_{i=1}^{n_1} \log t_{ij} \right) \right\}$$

$$p(\lambda_j | \dots) = Ga \left(\lambda_j | 2n_j + \lambda_0, \lambda_1 + \sum_{i=1}^{n_j} u_i \right) \quad j = 1, \dots, m$$

$$p(u_i | \dots) \propto \exp(-u_i \lambda) I(t_i^\alpha < \lambda) \quad i = 1, \dots, n$$

- 5 The distribution of a tessellation structure is given by

$$p(T | t_1, \dots, t_n, u_1, \dots, u_n) \propto p(t_1, \dots, t_n | u_1, \dots, u_n, T) p(u_1, \dots, u_n | T) p(T)$$

Thus we require the marginal distribution

$$p(t_1, \dots, t_n, u_1, \dots, u_n) = p(t_1, \dots, t_n | u_1, \dots, u_n) p(u_1, \dots, u_n)$$

- 10 The first term on the right hand side is given by

$$p(t_1, \dots, t_n | u_1, \dots, u_n) = \int_a^b \prod_{i=1}^n p(t_i | u_i, \alpha) p(\alpha) d\alpha$$

$$= \left(\prod_{i=1}^n u_i t_i \right) \int_a^b \alpha^{n + \alpha_0 - 1} \exp \left(\alpha \left(\sum_{i=1}^n \log t_i - \alpha_1 \right) \right) d\alpha$$

If $m = n + \alpha_0 - 1$ is an integer this integral can be evaluated by parts as follows

$$I_m = \int_a^b x^m \exp(xs) dx$$

$$= \left[x^m \frac{\exp(xs)}{s} \right]_a^b - \frac{m}{s} I_{m-1}$$

$$= \frac{1}{s} \sum_{i=0}^m \left[x^{n-i} \left(-\frac{n}{s} \right)^i \exp(xs) \right]_a^b$$

The marginal distribution of the latent variables is given by

$$\begin{aligned}
 p(u_1, \dots, u_n) &= \int_a^b \prod_{i=1}^n p(u_i | \lambda) p(\lambda) d\lambda \\
 &= \frac{\lambda_1^{\lambda_0} \prod_{i=1}^n u_{1i}}{\Gamma(\lambda_0)} \int_a^b \lambda^{2n+\lambda_0-1} \exp\left\{-\lambda \left(\lambda_1 + \sum_{i=1}^n u_i\right)\right\} d\lambda \\
 &= \prod_{i=1}^n u_{1i} \frac{\lambda_1^{\lambda_0} \Gamma(2n + \lambda_0)}{\Gamma(\lambda_0) \left(\lambda_1 + \sum_{i=1}^n u_i\right)^{2n+\lambda_0}}
 \end{aligned}$$

Given the marginal distribution $p(t_1, \dots, t_n, u_1, \dots, u_n) = p(t_1, \dots, t_n | u_1, \dots, u_n) p(u_1, \dots, u_n)$ the

- 5 tessellation structure is sampled for using a Metropolis random walk within the Gibbs (or other) sampler.

A range of applications are within the scope of the invention. These include situations in which it is required to predict whether a specified event will occur for an entity after a specified trigger event has occurred for that entity. For example, to if and when a customer
 10 will leave a bank after that customer has closed a loan with the bank. Other examples include predicting the lifetime of a patient after that patient has contracted a particular disease.

References

- Stephen G Walker and Eduardo Guterrez-Pera "Robustifying Bayesian Procedures"
 15 University of Valencia, Sixth Valencia International meeting on Bayesian Statistics, Invited papers, May 30 to June 4 1998.

Chen, Ibrahim and Sinha "A new Bayesian Model For Survival Data With a Surviving Fraction" Journal of the American Statistical Association, 1999.